

APPLICATION FOR UNITED STATES LETTERS PATENT

For

**ATCA INTEGRATED HEATSINK AND CORE POWER DISTRIBUTION
MECHANISM**

Inventors:

Edoardo Campini
William F. Handley
Mark D. Summers
Javier Leija

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard
Los Angeles, CA 90025-1026
(206) 292-8600

Attorney's Docket No.: 42P19177

"Express Mail" mailing label number: EV320117704US

Date of Deposit: March 31, 2004

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service

"Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to Mail Stop Patent Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Luci M. Arevalo

(Typed or printed name of person mailing paper or fee)

(Signature of person mailing paper or fee)

(DATE SIGNED)

ATCA INTEGRATED HEATSINK AND CORE POWER DISTRIBUTION MECHANISM

FIELD OF THE INVENTION

[0001] The field of invention relates generally to computer and
5 telecommunications equipment, and, more specifically but not exclusively relates to
an integrated heatsink and core power distribution mechanisms for computer and
telecommunication equipment boards.

BACKGROUND INFORMATION

10 **[0002]** During the 1990's the telecommunications market experienced two
extremely favorable market factors of quickly increasing bandwidth requirements
combined with fast market growth. While these factors were great for the industry,
they led to a hodgepodge of proprietary standards designed to meet the immediate
market needs that each company was working to service. Despite the advent of the
15 highly successful Compact PCI (peripheral component interconnect) standard, the
fact remains that the majority of the telecommunications market has no standard
form factor, backplane or fabric interconnect that can meet the 10Gb/sec+ bandwidth
requirements of today. There continues to be tens, if not hundreds, of different
chassis designs in the telecommunication industry, which drives the cost of this
20 equipment higher, while preventing multi-sourcing and interoperability across a
common backplane or fabric.

[0003] Significant changes in the telecommunications market have occurred in
the last few years, resulting in a leveling off of the meteoric growth of the 1990's,
with competition increasing accordingly. In today's cost-cutting environment of
25 constrained budgets, it is more important than ever for telecommunications
equipment providers to leverage off-the-shelf components and sub-systems, thereby
minimizing investment and maximizing the breadth of product environment.

[0004] In view of the foregoing, as well as other considerations, various industry efforts have been directed to defining a standardized, modular telecommunications equipment solution. A leading standard resulting from one of these endeavors is the Advanced Telecommunications Architecture (ATCA) (also referred to as Advanced TCA). The ATCA standard defines an open switch fabric based platform delivering an industry standard high performance, fault tolerant, and scalable solution for next generation telecommunications and data center equipment. The development of the ATCA standard is being carried out within the PCI Industrial Computer Manufacturers Group (PICMG) - the same group that created the highly successful Compact PCI standard. The Advanced TCA 3.0 base specification defines the physical and electrical characteristics of an off-the-shelf, modular chassis based on switch fabric connections between hot-swappable blades. The Advanced TCA base specification supports multiple fabric connections, including the advanced switching (AS) standard.

[0005] The ATCA 3.0 base specification defines the frame (rack) and shelf (chassis) form factors, core backplane fabric connectivity, power, cooling, management interfaces, and the electromechanical specification of the ATCA-compliant boards. The electromechanical specification is based on the existing IEC60297 EuroCard form factor, and enables equipment from different vendors to be incorporated in a modular fashion and be guaranteed to operate.

[0006] The ATCA 3.0 base specification also defines a power budget of 200 Watts (W) per board, enabling high performance servers with multi-processor architectures and multi gigabytes of on-board memory. The framepower is delivered by redundant -48 VDC feeds. These dual framepower feeds are typically fused and multiple sub-feeds are generated, allowing each shelf to remain electrically isolated. Local DC-DC conversion is accomplished per board. Redundant local power feeds

are normally attached through either diode *Or'ed* connections to a single on-board DC-DC converter or via on-board dual redundant load sharing DC-DC converters.

[0007] The PICMG organization has performed extensive thermal modeling in order to design the ATCA board and shelf form factors to be able to support 200W power dissipation per board slot. In a typical implementation, the shelf uses conventional air cooling, with blowers pulling air from front to rear and bottom to top. Mechanical fans are typically the element with the lowest MTBF (mean time between failure), and thus thermal designs should incorporate sufficient overhead to accommodate a failed blower. Blowers may be synchronized and include temperature controlled airflow to reduce audible levels.

[0008] Although the ATCA standard is a step in the right direction, several of its design aspects are limiting. Notably, the 200W power dissipation per board slot places a restriction on how much heat can be generated by a board, while not defining the location of the heat sources. (It is noted the 200W limit applies to boards occupying a single slot. Higher power dissipations are allowed for boards occupying multiple slots, e.g., a front board occupying two slots is limited to 400W). In order to meet ever increasing bandwidth demand, equipment capabilities need to increase, requiring faster processors, among other board component improvements. When considering identical semiconductor fabrication processes, there is a direct relationship between speed and power consumption – the faster the processor speed, the greater the power consumption of the processor, and thus the greater the amount of heat that must be dissipated by an ATCA board hosting the processor. In view of the rigid board and shelf form factors defined by the ATCA base specification, this leads to problems in achieving sufficient cooling for high-power components such as processors, since the majority of the 200W will be consumed by such components. Power routing on the boards is also problematic.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The foregoing aspects and many of the attendant advantages of this invention will become more readily appreciated as the same becomes better understood by reference to the following detailed description, when taken in conjunction with the accompanying drawings, wherein like reference numerals refer to like parts throughout the various views unless otherwise specified:

[0010] Figure 1 is a schematic diagram of an exemplary airflow pattern defined by the ATCA (Advanced Telecommunication Architecture) standard for cooling ATCA boards;

[0011] Figure 2a is an isometric view of an ATCA Front Board including first and second power rails that are employed to provide power to first and second processors while simultaneously providing a mounting mechanism to thermally couple first and second heatsinks to the first and second processors, according to one embodiment of the invention;

[0012] Figure 2b is an isometric view of the ATCA Front Board embodiment of Figure 2a showing the configuration of the heatsinks after assembly;

[0013] Figure 2c is an isometric view showing further details of power pads, power supply pads, and IC (integrated circuit) pads formed in the printed circuit board (PCB) of the ATCA Front Board of Figure 2a, according to one embodiment of the invention;

[0014] Figure 2d is an isometric view showing further details of an alternative configuration for power pads formed in the PCB of the ATCA Front Board of Figures 2a-b, according to one embodiment of the invention;

[0015] Figure 2e is an isometric view showing a variation of the configuration of the ATCA Front Board embodiment of Figures 2a-b, wherein an additional heat sink is disposed toward the top of the first and second power rails;

[0016] Figure 2f is an isometric view showing a variation of the configuration of the embodiment of Figures 2a-b, wherein a single elongated heat sink is used;

[0017] Figure 3a is a cross-sectional view taken along line 3a-3a of Figure 2b, illustrating further details of the first and second power rails, according to one
5 embodiment of the invention; and

[0018] Figure 3b is a cross-sectional view taken along line 3a-3a of Figure 2b, illustrating further details of first and second power rails having a split-rail configuration, according to one embodiment of the invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0019] Embodiments of apparatus and methods employing integrated heatsink and core power distribution mechanisms are described herein. In the following description, numerous specific details are set forth, such as implementations
5 corresponding to the ATCA standard, to provide a thorough understanding of embodiments of the invention. One skilled in the relevant art will recognize, however, that the invention can be practiced without one or more of the specific details, or with other methods, components, materials, *etc.* In other instances, well-known structures, materials, or operations are not shown or described in detail to
10 avoid obscuring aspects of the invention.

[0020] Reference throughout this specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, the appearances of the phrases “in one embodiment” or “in
15 an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable manner in one or more embodiments.

[0021] In the following paragraphs, exemplary embodiments of the invention are
20 described within a use context targeted to ATCA equipment installations. However, this does not limit the use of the principles and teachings disclosed herein to ATCA equipment. Rather, the following embodiments are merely illustrative of one possible type of installation. In general, the principles and teachings are applicable to various types of modular electronic equipment, including, but not limited to,
25 telecommunications equipment and computer equipment.

[0022] The majority of equipment installed in central office and data centers today is air cooled due to simplicity, relatively low cost, ease of implementation, and

reliability. This trend is likely to continue into the future for the same reasons. Accordingly, the thermal guidelines defined in the ATCA specification apply to air cooling; other cooling methods are permissible, but are not covered in the specification. In general, Boards, Shelves, and Frames *may* be either cooled by
5 natural convection, without the assistance of fans or blowers, or cooled by forced convection with the assistance of fans or blowers. The choice is left to the end-user requirements.

[0023] An airflow path used as an example in the ACTA PICMG 3.0 Short Form Specification (January, 2003) is shown in Figure 1. Air enters a Shelf 100 at in air
10 inlet 102 at the lower front into a bottom plenum 104 and turns 90 degrees upward. Airflow across the bottom edge of a Front Board 106 and a rear transition module (RTM) 108 is evenly distributed. As the air passes across the hot components on Front Board 106 and RTM 108, heat is carried away by forced convection. The air
15 exits the subrack at the top, is drawn into an upper plenum 110, turns 90 degrees, and is exhausted out the rear of Shelf 100 by fans 112. In addition to the illustrated configuration, other airflow paths and cooling methods are allowed under the ACTA standard.

[0024] In accordance with aspects of the embodiments described herein, enhanced cooling is facilitated through use of a pair of power rails, which are used to
20 assist in removing heat from high power components, such as processors and the like. Furthermore, the power rails also function as buss bars to supply low voltage, high current electrical power to the high-power components being cooled.

[0025] One exemplary embodiment of this scheme is shown in Figures 2a, 2b, and 3a. Each of Figures 2a-f show a Front Board 200 comprising a printed circuit
25 board 202 having a front panel 204 coupled thereto. Various components are coupled to or otherwise mounted on PCB 202, including a power module 206 disposed toward the top of PCB 202 and connectors 208, 210, and 212.

[0026] In general, the power module 206 includes one or more DC-DC converters, which convert input power (typically -48 V DC) to appropriate voltage levels to supply power at those voltage levels to various selected board components. (The internal details of the power module are not shown herein for clarity –
5 techniques for designing and implementing power supplies of this type are well-known in the art; accordingly, further details are not discussed herein). These components include a pair of processors 214A and 214B. Other components powered by power module 206 are not shown here for clarity. In general, processors 214A and 214B may be coupled to Front Board 200 using one of several
10 conventional mounting schemes, including soldering directly to a set of pads on a top layer of Front Board 200 or using a socket-type connections, wherein a mating socket is mounted to a set of drill-through pads (typically).

[0027] The embodiment of Figures 2a-b and 3a is designed to be used for any conventional type of integrated circuit (IC), such as, but not limited to, a
15 microprocessor, a network processor, an application specific integrated circuit (ASIC), a field-programmable gate array (FPGA), a microcontroller, etc. Furthermore, the IC's packaging includes, but is not limited to, a Pin Grid Array (PGA) and a Ball Grid Array (BGA). As such, each of these types of IC's includes one or more pins (for PGA's) or solder pads (for BGA) to which power is supplied to
20 make the IC operate.

[0028] Under a conventional PCB configuration, power is usually supplied via power traces defined in one or more layers of the PCB. In some instances, an entire layer is (substantially) used for power routing (e.g., for routing a primary positive voltage or ground). These layers are referred to as power planes and ground
25 planes, respectively. Under other common configurations, traces supplying positive voltages and functioning as grounds may be provided within the same layer, while being separated by an appropriate insulator, such as the PCB substrate material.

The power traces eventually are coupled to specific pads that are used to either directly connect to power pins or solder balls on the IC, or to socket pins that are internally connected to such power pins. This is done through various routing layers and vias, as is well-known in the PCB art.

5 **[0029]** One disadvantage with using power planes for high-power boards is that the amount of area occupied by the power traces tends to get large, increasing routing complexities for other board components and increasing heat due to resistive electrical losses. This problem is especially noticed when supplying power to low voltage, high current components, such as modern processors. For example, the
10 core voltage of many of today's processors is only 2.0 volts or less (*e.g.*, 1.0 – 1.5 volts). Thus, at 1.0 volts, a processor that is rated at 100 Watts would consume 100 amps of current when operating at its peak power rating. When multiple processors of this type are employed, the problem becomes even more pronounced.

15 **[0030]** In the embodiments of Figures 2a-f, as well as other embodiments described herein, the foregoing power routing problem is address by employing power "rails" that function as buss bars. The power rails have a relatively large cross-section (when compared with conventional power traces or power planes), leading to a dramatic reduction in heat generated from resistive losses. Furthermore, the power rails provided a convenient means for providing power to
20 various board components, and significantly reduce routing real estate.

25 **[0031]** As shown in Figures 2a and 2b, first and second power rails 216A and 216B are used to supply power to processors 214A and 214B, wherein the power rails function as buss bars. In essence, each power rail performs a similar function to a power trace defined in a layer of PCB 202, except with much lower resistance due to the dramatically-increased cross section of a power rail when compare with a PCB trace. Accordingly, there needs to be a mechanism for electrically coupling one or more pads corresponding to pins (or solder balls) on an IC to which power at a

give voltage is to be provided with by a corresponding power rail (or portion of a split power rails, as described below in further detail) with that power rail.

[0032] An exemplary power, IC, and power supply pad layout according to one embodiment is shown in Figure 2c. As is well-known, pads are used to mechanically and electrically coupled surface-mounted components to PCBs by using an appropriate soldering technique. Drill-through pads may be used for coupling components with "pinned" packaging, while non-drilled pads are used for components that use non-pinned packaging, such as solder-ball coupling techniques. For illustrative purposes, all of the pads shown in Figures 2c and 2d are drill-through pads. However, any type of pad may be used, depending on the corresponding component that is to be mounted to the PCB.

[0033] Accordingly, Figures 2c and 2d show various pads used to couple corresponding components to PCB 202, including a set of power supply pads 218 and sets of IC pads 220A and 220B. Power supply pads are illustrative of pads used to couple power supply components to PCB 202. In general, power module 206 is illustrative of a power supply used to supply power to the various components of Front Board 200. In addition to the monolithic power brick configuration illustrated in Figures 2a and 2b, the power supply may comprise individual components mounted or otherwise coupled to PCB 202. These components will typically include DC-DC converters and various power-conditioning circuitry. The power supply circuitry will also include various forms of fusing and/or employ other isolation and circuit protection mechanisms to prevent damage to Front Board 200.

[0034] Under the PICMG 3.0 specification, each board is supplied with dual -48 VDC feeds. Each feed remains isolated and is fed individually to each board slot through the backplane. There are two basic methods for combining the dual redundant feeds. One method combines the two feeds through diode OR'ing and

delivers the combined single feed to DC-DC converters. If either feed fails, all power shall be delivered by the surviving feed. The second method is to direct each of the two feeds to its own DC-DC converter. The outputs of the converters are then combined to provide power to the rest of the power supply circuits.

5 **[0035]** Under some embodiments, power module 206 is illustrative of a power supply comprising a field replaceable unit (FRU). Accordingly, an appropriate connector or set of connectors is/are coupled to PCB 202 to enable power module 206 to electrically couple to the PCB via a mating set of power supply pads, such as depicted by power supply pads 218.

10 **[0036]** IC pads 220A and 220B are illustrative of pads used to couple processors 214A and 214B to PCB 202. These sets of pads are used to couple electrical power and signals to an IC that is either directly coupled to the set of IC pads or is coupled to a socket connector or the like that is directly coupled to the set of IC pads.

15 **[0037]** In order to supply power to processors 214A and 214B, there needs to be a mechanism for routing power to selected pads from among IC pads 220A and 220B. In the embodiment of Figure 2c, this is facilitated by power pads 222, 224, 226, 228, 230, and 232. As illustrated, power pads 220 and 224 are used to connect selected power supply pads 218 to power rails 216A and 216B, respectively (shown
20 in hidden lines). Also as shown, each of power pads 220 and 224 are formed in a top layer of PCB 202. Optionally, a portion of these power pads may be routed via an internal layer in the PCB. Also, a non-conductive coating or the like may be applied over the portion of power pads 222 and 224 that is under power module 206.

25 **[0038]** Power pads 226 and 228 are used to electrically couple power rails 216A and 216B to selected pads from among IC pads 220B. For example, in one embodiment, power pad 226 is connected to a set of supply voltage IC pads for processor 214B, while power pad 228 is connected to a set of ground IC pads for

processor 214B. Power pads 230 and 232 are used to electrically couple power rails 216A and 216B to selected pads from among IC pads 220A in a similar manner.

5 **[0039]** In the power pad routing schemes of Figures 2c and 2d, power pads 226, 228, 230, and 232 are shown on the top layer of PCB 202. However, this is merely one exemplary configuration. In general, power pad routing may employ use of internal layers, as well as appropriate vias, as is well-understood in the art.

10 **[0040]** In order to route power via power rails 216A and 216B, there need to be a mechanism for adequately coupling the power rails to power pads 222, 224, 226, 228, 230, and 232. In one embodiment, power rails 216A and 216B are coupled to PCB 202 using various types of fasteners, such as machine screws, shoulder screws, etc, wherein the shanks of the fasteners pass through the PCB and thread into the power rails. An exemplary mounting scheme of this type, using fasteners 340, is shown in Figure 3a. In another embodiment, solder is used to
15 electrically (and mechanically) couple the power rails to the PCB using an appropriate soldering technique, such as wave-soldering. As yet another option, a power rail may be coupled to the power pad(s) using a conductive epoxy. In still another embodiment, a combination of the foregoing schemes may be employed, such as using a combination of fasteners and soldering, for example.

20 **[0041]** In one embodiment, a power pad having a configuration similar to the planer profile of a power rail is formed on the top layer of PCB 202, such as shown in Figure 2d. This configuration is analogous to that shown in Figure 2c, except for the addition of elongated power pads 234 and 236. It is recommended that solder not be applied along the entire length of elongated power pads 234 and 236, since
25 there is a coefficient of expansion mismatch between the layer substrate of PCB 202 (e.g., fiberglass) and the material used for power rails 214A and 214B, which are typically made of a metal with good conduction, such as aluminum, copper, or

various well-known alloys. This coefficient of expansion mismatch may lead to fatigue failure of the solder joints.

[0042] Another aspect of the power rails employed in the embodiment of Figures 2a-f, and 3a-b is the ability to function as a heat sink rail. For example, in the embodiment of Figure 3a, each of power rails 216A and 216B has a profile including one or more slots 300_n (e.g., 300₁ and 300₂ for a two-slot configuration) defined in an inward-facing edge of the profile. In the illustrated embodiment of Figure 3a, an optional second set of slots 301₁ and 301₂ are formed in the outward facing edge of power rails 216A and 216B, as well. These second sets of slots are used to enhance cooling of the power rails via forced or natural convection.

[0043] The slots in the power rails are configured to slidably engage a corresponding flange 302 defined in a heatsink 304. As shown in the assembly process of Figures 2a and 2b, as each of heatsinks 304A and 304B is placed over a respective processor 214A and 214B, the flanges of the heatsinks slide within respective slots 300_n in power rails 216A and 216B. The slots to be employed are dependent on the height of the top surface of an installed processor. Thus, a power rail with multiple slots provides support for processor installations having multiple different heights.

[0044] In general, it will be desired to securely hold heatsinks 304A and 304B in place upon installation, while thermally coupling the heatsinks to their respective processor. This may be performed by one of several different schemes, such as, but not limited to, employing fasteners, wedge blocks, spring clips, etc. The securing mechanism may also be used to adjust the height of the base of a heatsink relative to the top surface of a processor. At the same time, it is desired to thermally couple heatsink 304 to either (substantially) all of the upper surface of processor 214, or at least the area proximate to the die area of the processor. This can be done using several techniques, including employing a thermal paste 306.

[0045] Another consideration relates to electrical isolation of the power rails. In some embodiments, power having different voltage levels is supplied to power rails 216A and 216B, such as processor core power (e.g., 2.0 volts) supplied to one power rail while employing the other power rails as a ground. Since heatsinks are
5 typically made of conductive materials, such as aluminum, mechanically coupling a heatsink across the two power rails will cause a short. This is highly undesirable.

[0046] This consideration may be addressed in one of several fashions. In the embodiment illustrated in Figure 3a, a thermally conductive, yet non-electrically-conductive, coating 308 is applied to one or both of slots 300n and heatsink
10 flanges 302. For example, there are various coatings that are both thermally conductive and good electrical insulators, such as thermally-conductive paints and ceramics.

[0047] In accordance with another technique for enhancing the cooling of processors 214A and 214B, either or both of power rails 216A and 216B may
15 employ a heat pipe 310. A typical heat pipe consists of a sealed metal container housing a capillary wicking material and a working fluid. A heat pipe provides the ability to transport heat against gravity by employing an evaporation-condensation cycle with the help of porous capillaries that form the wick. The wick provides the capillary driving force to return the condensate to the evaporator section of the heat
20 pipe. The quality and type of wick usually determines the performance of the heat pipe. Different types of wicks are used depending on the application for which the heat pipe is being used.

[0048] The working fluid is used to transfer heat via fluid transport in combination with an evaporation/condensation cycle. Thus, a first consideration in the
25 identification of a suitable working fluid is the operating vapor temperature range. Within the approximate temperature band, several possible working fluids may exist, and a variety of characteristics should be examined in order to determine the most

acceptable of these fluids for the application considered. The selection of the working fluid should also be based on thermodynamic considerations which are concerned with the various limitations to heat flow occurring within the heat pipe, such as viscous, sonic, capillary, entrainment and nucleate boiling levels.

5 **[0049]** In heat pipe design, a high value of surface tension is desirable in order to enable the heat pipe to operate against gravity and to generate a high capillary driving force. In addition to high surface tension, it is necessary for the working fluid to wet the wick and the container material - *i.e.*, the contact angle should be zero or very small. The vapor pressure over the operating temperature range must be
10 sufficiently great to avoid high vapor velocities, which tend to set up large temperature gradients and cause flow instabilities.

[0050] A high latent heat of vaporization is desirable in order to transfer large amounts of heat with minimum fluid flow, and hence to maintain low pressure drops within the heat pipe. The thermal conductivity of the working fluid should preferably
15 be high in order to minimize the radial temperature gradient and to reduce the possibility of nucleate boiling at the wick or wall surface. The resistance to fluid flow will be minimized by choosing fluids with low values of vapor and liquid viscosities.

[0051] The prime purpose of the wick is to generate capillary pressure to transport the working fluid from the condenser to the evaporator. It must also be
20 able to distribute the liquid around the evaporator section to any area where heat is likely to be received by the heat pipe. Often these two functions require wicks of different forms. The selection of the wick for a heat pipe depends on many factors, several of which are closely linked to the properties of the working fluid.

[0052] The wicking material is a porous structure made of materials such as
25 steel, aluminum, nickel or copper in various ranges of pore sizes. These structures are typically fabricated using metal foams and felts (most common). By varying the pressure on the felt during assembly, various pore sizes can be produced. By

incorporating removable metal mandrels, an arterial structure can also be molded in the felt. The most common types of wicks are sintered powders, grooved tubes, and screen meshes.

[0053] Fibrous materials, like ceramics, may also be used. These materials generally have smaller pores. The main disadvantage of ceramic fibers is that they have little stiffness and usually require a continuous support by a metal mesh. Thus while the fiber itself may be chemically compatible with the working fluids, the supporting materials may cause problems. More recently, interest has turned to carbon fibers as a wick material. Carbon fiber filaments have many fine longitudinal grooves on their surface, have high capillary pressures and are chemically stable.

[0054] The maximum capillary head generated by a wick increases with decrease in pore size. The wick permeability increases with increasing pore size. Another feature of the wick, which must be optimized, is its thickness. The heat transport capability of the heat pipe is raised by increasing the wick thickness. The overall thermal resistance at the evaporator also depends on the conductivity of the working fluid in the wick. Other necessary properties of the wick are compatibility with the working fluid and wettability.

[0055] A heat pipe operates in the following manner. Inside the sealed container is a liquid under its own pressure, which enters the pores of the capillary material, wetting all internal surfaces. Applying heat at any point along the surface of the heat pipe causes liquid working fluid at that point to boil and enter a vapor state. When that happens, the working fluid picks up the latent heat of vaporization. The gas, which then has a higher pressure, moves inside the sealed container to a colder location where it condenses. Thus, the gas gives up the latent heat of vaporization and moves heat from the input (at which a heat source is located) to the output end of the heat pipe. In the embodiments of Figures 2a-f, the output end is disposed toward the top of the Figures.

[0056] Returning to Figure 3a, supporting a heat pipe in a power rail can be enabled in the following manner. First, a cavity is formed in a power rail. In the illustrated embodiment of Figure 3a, the cavity has a circular cross-section, which may be formed using one of several techniques, such as extruding a power rail.
5 Optionally, the cavity may be formed by machining or casting. Furthermore, the internal profile of the cavity is generally not important.

[0057] Next, an appropriate wicking material is inserted into the cavity. As discussed above, various types of wicking material may be used. In one embodiment, a sintered metal cast in the shape of the power rail cavity is used.
10 cases in which an extruded profile is used, each end of the profile needs to be sealed. When machining or casting is employed, only one end of the profile may need to be sealed (the other is inherently sealed by the power rail material) in some instances. Generally, it will be preferred that one end be sealed, while the other remains unsealed before inserting the wicking material.

15 **[0058]** Either before or after the wicking material is inserted into the cavity, the working fluid is added. It is important that the working fluid not completely fill the cavity, as this will prevent proper operation of the evaporation/condensation cycle. Additionally, the working fluid should be selected based on the expected operating temperature range for the heat pipe. After the wicking material and working fluid is
20 added, the cavity is sealed to enclose the working fluid and wicking material. Typically, this can be accomplished by covering an open end of the power rail with a thin material, such as a metal. Various types of sealing techniques may be employed for this purpose, such as soldering, brazing, using epoxy, *etc.* The net objective is to hermetically seal both ends of the power rail, thus preventing the
25 working fluid from escaping.

[0059] In addition to supplying a single voltage level to various circuitry, an individual power rail may be employed to supply multiple voltage levels. For

example, a split-rail embodiment for a power rail is shown in Figure 3b. Under this embodiment, one or both of power rails 216A and 216B (both shown here) have a "split" configuration. The split terminology refers to two conductive sections 312 and 314 being divided (split) by an insulator section 316. These three sections may be
5 joined by using a suitable assembly technique. In one embodiment, conductive sections 312 and 314 are simply joined using a suitable non-conductive adhesive comprising insulator section 316.

[0060] In addition to the two conductive sections shown in Figure 3b, additional conductive sections may be employed. Under such embodiments, each adjacent
10 pair of conductive sections will be separated by an insulator section. It is noted that the material used for an insulator section does not have to occupy the entire gap between the two conductive sections. For example, a portion of the gap may be occupied by air.

[0061] In general, power from the different conduction sections will be routed to
15 selected IC pads 220 via different layers in PCB 202 (although the same layer may be used). Under these circumstances, appropriate power traces and vias will be employed. For illustrative purposes, an exemplary routing scheme employing power pads 318, 320, 322, and 324 along with power traces 326 and 328 and vias 330, 332, 334, and 336 is shown in Figure 3b.

[0062] In one embodiment, the supply voltages (or ground) supplied to the inner
20 conductive sections 312 are the same. Accordingly, upon energizing power rails, there is no voltage differential across the heatsinks 304A and 304B. Thus, there is no need for thermally conductive coating 308 in this embodiment.

[0063] As shown in Figures 2e and 2f, the heatsinks used in an actual
25 implementation may extend over more than just corresponding high-power ICs, such as processors 214A and 214B. For example, the scheme of Figure 2e employs a heatsink 304C disposed toward the top ends of power rails 216A and 216B. A

heatsink in this location may be advantageous for use in embodiments employing heatpipes. Under the scheme of Figure 2f, a single elongated heatsink 304D runs substantially the length of power rails 216A and 216B.

5 **[0064]** Using computer-based heat transfer modeling, embodiments configured in accordance with the principles and techniques disclosed herein have been shown to dramatically enhance cooling. For example, one model shows a 40% reduction in package temperature for the embodiment of Figures 2a and 2b, when compared with a similar configuration that includes heatsinks analogous to heatsinks 304A and 304B but does not employ power rails 216A and 216B.

10 **[0065]** As discussed above, using power rails also saves significant routing real estate. For example, using a conventional PCB structure requires a power trace width of approximately 1.5 inches to supply 100A of current (the amount of current corresponding to a 200W consumption by components running at 2.0 V core voltages). Thus, since power traces of this size are required both for the supply
15 voltage and ground, approximately 3 inches of board real estate, running the entire length of a board, would be required. In contrast, in one embodiment power rails 216A and 216B are approximately .25 inches in width, thus requiring about one-half inch of board real estate for two rails.

20 **[0066]** In addition to the foregoing power distribution and cooling functions, power rails also serve a board-stiffening function. The schemes disclosed herein enable a designer to randomly locate one or more heatsinks to cool one or more components of varying heights. Furthermore, the power rail configuration also minimizes the number of holes in a PCB required to fasten heatsinks.

25 **[0067]** The above description of illustrated embodiments of the invention, including what is described in the Abstract, is not intended to be exhaustive or to limit the invention to the precise forms disclosed. While specific embodiments of, and examples for, the invention are described herein for illustrative purposes,

various equivalent modifications are possible within the scope of the invention, as those skilled in the relevant art will recognize.

[0068] These modifications can be made to the invention in light of the above detailed description. The terms used in the following claims should not be construed
5 to limit the invention to the specific embodiments disclosed in the specification and the claims. Rather, the scope of the invention is to be determined entirely by the following claims, which are to be construed in accordance with established doctrines of claim interpretation.